

Course Information

CS 59200 – Probabilistic Machine Learning

Semester: Fall 2022

Lecture: TTh 12:00 - 1:15p @ LWSN B134

Instructor: Ruqi Zhang

Email: ruqiz@purdue.edu

Prerequisites: Knowledge of basic probability and machine learning

Course Description

How can a machine learn from data? How should we construct appropriate models? How do we design algorithms for inference? What are the fundamental design principles of data-driven learning and decision-making?

This course will help you answer these fundamental questions in machine learning from a probabilistic perspective. You will read and discuss research papers in probabilistic machine learning, including but not limited to research topics about

1. **Probabilistic models:** Bayesian neural networks, variational autoencoder, normalizing flows, score-based generative models, diffusion models, Gaussian processes.
2. **Probabilistic inference:** Markov chain Monte Carlo (MCMC), variational inference, Laplace approximation, dropout, ensemble.
3. **Evaluations and applications:** uncertainty estimation, calibration, distribution shift, out-of-distribution detection, adversarial attacks.

This course is designed to introduce techniques and research topics in probabilistic machine learning. Students at all levels are welcome. PhD, master and undergrad students who are interested in doing a research project in related fields are particularly encouraged to take this course.

Format

The class will be a mix of lectures and student presentations. Students will be mostly presenting and discussing research papers from top-tier ML venues, followed by breakout discussions about the material. I will give some lectures on probabilistic machine learning foundations.

Learning Outcomes

By the end of the semester, you should be able to:

1. Think about any ML problem from a probabilistic perspective.
2. Be familiar with common probabilistic models and inference methods.
3. Know the style of academic writing in ML.
4. Know how to critically assess a ML paper.
5. Formulate and carry out a short-term research project.

Learning Resources, Technology & Texts

There is no required textbook for the course. Some recommended textbooks:

- [*Probabilistic Machine Learning*](#) (a book series), Kevin Murphy
- [*Pattern Recognition and Machine Learning*](#), Christopher Bishop
- [*Information Theory, Inference, and Learning Algorithm*](#), David MacKay
- [*Deep Learning*](#), Ian Goodfellow, Yoshua Bengio and Aaron Courville

Grading

Reading (15%)

Paper presentation (25%)

Final project (55%)

Discussion participation (5%)